

Trends in Processor, Communications, and Connection Technologies

Martin G. Moebus
Surface Undersea Warfare Department



**Naval Undersea Warfare Center Division
Newport, Rhode Island**

Approved for public release; distribution is unlimited.

20011123 026

PREFACE

The research in this report was conducted under Task Instruction 5485, "AN/SQQ-89 Technology Refresh Project," principal investigator Martin G. Moebus (Code 3121). The sponsoring activity is the Program Executive Office, Mine and Undersea Warfare (PMS 411), program manager Eric Plummer (PMS 411D4).

The technical reviewer for this report was Geoffrey P. Zito (Code 3121).

The author gratefully acknowledges the contributions of Anish Shukla (University of Connecticut) for assistance in researching the AMD processor and other technology areas.

Reviewed and Approved: 28 September 2001



**Patricia J. Dean
Head, Surface Undersea Warfare Department**



REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 0704-0188*

Public reporting for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>
<p>1. AGENCY USE ONLY (Leave blank)</p> <p>2. REPORT DATE 28 September 2001</p> <p>3. REPORT TYPE AND DATES COVERED</p>			
<p>4. TITLE AND SUBTITLE Trends in Processor, Communications, and Connection Technologies</p>		<p>5. FUNDING NUMBERS E225008</p>	
<p>6. AUTHOR(S) Martin G. Moebus</p>			
<p>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Undersea Warfare Center Division 1176 Howell Street Newport, RI 02841-1708</p>		<p>8. PERFORMING ORGANIZATION REPORT NUMBER TD 11,313</p>	
<p>9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Naval Sea Systems Command (PMS 411D4) 614 Sicard Street Washington Navy Yard, DC 20376-7019</p>		<p>10. SPONSORING/MONITORING AGENCY REPORT NUMBER</p>	
<p>11. SUPPLEMENTARY NOTES</p>			
<p>12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.</p>		<p>12b. DISTRIBUTION CODE</p>	
<p>13. ABSTRACT (Maximum 200 words) Some of the underlying computing technologies that may be applied to future system development are examined. Current and soon-to-be-made available commercial industry technology is reviewed, including processors, communications, and connectivity.</p>			
<p>14. SUBJECT TERMS Computer Technology Processor Roadmap Technology Roadmap</p>			<p>15. NUMBER OF PAGES 38</p>
			<p>16. PRICE CODE</p>
<p>17. SECURITY CLASSIFICATION OF REPORT Unclassified</p>	<p>18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified</p>	<p>19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified</p>	<p>20. LIMITATION OF ABSTRACT SAR</p>

TABLE OF CONTENTS

	Page
LIST OF ILLUSTRATIONS.....	ii
LIST OF TABLES.....	iii
LIST OF ABBREVIATIONS AND ACRONYMS	iii
INTRODUCTION	1
TECHNOLOGY ROADMAPS	1
Computer Architecture.....	1
Product Life Cycles.....	3
PROCESSOR ROADMAPS	3
INTEL	5
IA-32 Architecture	5
IA-64 Architecture	5
Alpha.....	5
Sparc	6
PowerPC	6
Motorola ‘G’ Family.....	6
IBM	6
AMD	7
TYPES OF MEMORY	7
I/O INTERCONNECTION TECHNOLOGIES	7
Bus Architecture	8
PCI and PCI-X.....	8
VME.....	9
SCSI	10
Switched Fabric Connections	10
Fibre Channel.....	10
Race.....	11
RapidIO.....	11
InfiniBand	11
PLX	12
Starfabric.....	12
PERIPHERAL CONNECTIVITY.....	13
IEEE 1394.....	13
USB.....	13
Parallel Port/Serial Port.....	13
ISA Bus	13

TABLE OF CONTENTS (Cont'd)

	Page
NETWORKS	14
Ethernet Technology	14
Fast Ethernet	15
Gigabit Ethernet	15
10-Gigabit Ethernet (10GBASE-X)	15
Asynchronous Transfer Mode (ATM)	16
Sonet: High-Speed, Fiber-Based Transmission Medium	16
INTEGRATED PRODUCTS	17
STORAGE SOLUTIONS	17
Hard Drives	17
SCSI Drives	17
IDE Drives	17
Direct Attach Storage (DAS)	18
RAID	18
NAS	19
Storage Area Network (SAN)	19
SUMMARY	20
REFERENCES	21
APPENDIX—PROCESSOR ROADMAPS FOR VARIOUS VENDORS	A-1

LIST OF ILLUSTRATIONS

	Figure	Page
1	Highly Parallel System Architecture as Implemented in the Compaq Professional Workstation SP700	2
2	Technology Roadmap by Vendor	4
3	I/O Interconnection Technologies	8
4	PCI Plus Infiniband	12
5	Ethernet Technology: How Fast?	14
6	Example of RAID Drive	18
7	Example of a NAS Drive	19
8	SAN Architecture	20

LIST OF TABLES

Table		Page
A-1	Summary of Intel Processors	A-1
A-2	Summary of Alpha Processors.....	A-3
A-3	Summary of Sparc Processors	A-4
A-4	Summary of Motorola Processors.....	A-5
A-5	Summary of IBM Processors.....	A-6
A-6	Summary of AMD Processors	A-7

LIST OF ABBREVIATIONS AND ACRONYMS

AMD	Advanced Micro Devices
ATM	Asynchronous transfer mode
BISDN	Broadband integrated services digital networks
CMOS	Compatible metal oxide semiconductor
CPU	Central processing unit
CSMA/CD	Carrier sense multiple access/collision detection
DAS	Direct attach storage
DEC	Digital Equipment Corporation
DIMM	Dual inline memory module
DDR	Double-data rate
DRAM	Dynamic RAM
EIDE	Enhanced IDE
EPIC	Explicitly parallel instruction computing
EEPROM	Electronically erasable programmable read-only memory
FC	Fiber-optic connection
FDDI	Fiber-distributed data interface
ICMG	Industrial Computers Manufacturers Group
IDE	Integrated drive electronics
IP	Internet protocol
LAN	Local area network
MAC	Media access control
MAN	Metropolitan area network
NAS	Network attached storage
NIC	Network interface card

LIST OF ABBREVIATIONS AND ACRONYMS (Cont'd)

OC	Optical carrier
PCI	Peripheral component interconnect
QOS	Quality of service
RAID	Redundant array of independent disks
RDRAM	Rambus dynamic RAM
SAN	Storage area network
SCSI	Small computer system interface
SDRAM	Synchronous dynamic RAM
SIMM	Single inline memory module
SMP	Symmetric multiprocessing
SONET	Synchronous optical network
SRAM	Static RAM
TCP	Transmission control protocol
USB	Universal serial bus
UTP	Unshielded twisted pair
VME	Versa module Eurocard
WAN	Wide-area network

TRENDS IN PROCESSOR, COMMUNICATIONS, AND CONNECTION TECHNOLOGIES

INTRODUCTION

This report looks at some of the underlying computing technology that may be applied to future system applications. Its primary objective is to survey current and soon-to-be-on-the-market commercial computing technology. Roadmaps for processors, communication, and connection technologies are provided. Trends in graphics and displays are not included in this report.

TECHNOLOGY ROADMAPS

In order to evaluate total system performance, it is important to be aware of where technology is heading for all components of a computing system. Processor roadmaps are tied to the supporting chip sets that make up the processing system. Motherboards, or processor boards, create the fundamental processor platform used to build desired computing systems. Performance will also be linked to the central processing unit (CPU), along with other architectural issues, such as system processor bus speed, cache size and levels provided, random access memory (RAM) speed, and external bus speeds used to get information in and out of the processor. Graphic processors affect display performance. External storage devices and connection methods may also impact performance. Network architecture has an impact on processor communication speeds, and, thus, affects application performance. Focusing on a single parameter, such as processor speed, without considering the total computing architecture may lead to non-optimal choices in technology selection.

Commercial roadmaps look only several years ahead. Monitoring of product lifetimes is important when applying commercial technology to applications requiring long-term support. Constant vigilance in observing commercial trends and their applications to future system needs is essential for choosing system solutions that are maintainable in the long run.

COMPUTER ARCHITECTURE

Figure 1 shows the components of a typical computer system (reference 1). These include the CPU processor, memory and memory controllers, a graphics card, and I/O controllers (i.e., peripheral component interconnect (PCI) controllers). The I/O controllers allow connection of external or in-the-box peripherals, such as small computer system interface (SCSI) disks and network connections.

Until recently, most systems contained a single processor—hence, the term central processing unit, or CPU. The CPU executes the programs. Systems with only one processor are called serial, or scalar, processors.

A separate graphics processor is used to assist in rendering graphics data. This graphics processor interfaces directly to a memory controller by a separate graphics bus.

A memory controller links processors, graphics, memory, and I/O functions together, and allows access and retrieval to onboard RAM. Instructions and data needed by the processor are stored in the RAM while waiting for and returning from CPU processing. The RAM is stored in dual inline memory modules (DIMMs) or in single inline memory module (SIMM) packages. Programs and data are generally stored on a hard drive for long-term storage.

The system bus, or front side/processor bus, allows the CPU, memory, I/O controllers, and graphics processors to communicate. I/O controllers, like a PCI, use a separate bus (a PCI bus) to communicate externally in the box and to the outside world.

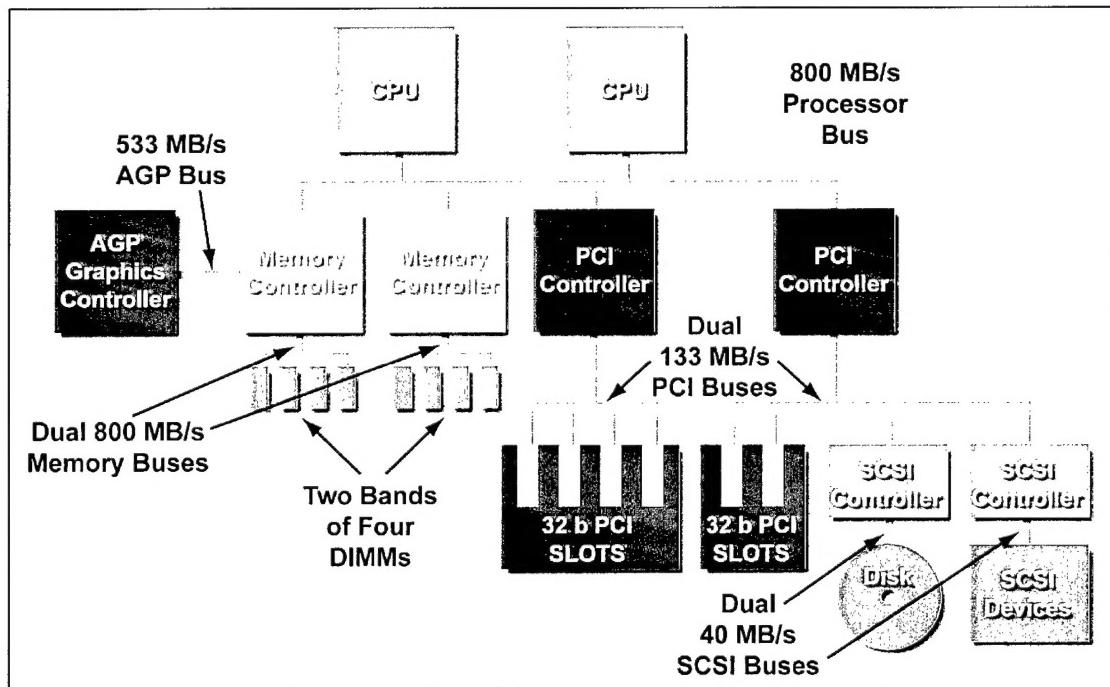


Figure 1. Highly Parallel System Architecture as Implemented in the Compaq Professional Workstation SP700

Many processors use caches to allow transactions to proceed faster by storing data and instructions closer to the CPU. This small memory device runs at a faster speed than RAM memory. Some vendors have up to three levels of backside cache. Level 1 (L1) is generally built into the CPU and runs at the same speed as the processor; level 2 cache (L2) runs a little slower and may also be on the CPU chip. The level 3 cache (L3) is generally on a separate chip.

The I/O controller allows for data and programs to be sent to and from memory to external storage devices, such as a disk drive or tape device. These external devices operate at a much slower rate than internal memory. Other I/O devices include operator-machine interface devices, such as a display, keyboard, and mouse. Another important connection device is the network interface card (NIC). A NIC allows connection to a network. The NIC works at the physical layer (i.e., physical connection) of the standard seven-layer OSI networking model. The NIC conveys the bit stream (electrical, optical, etc.) through the network at the electrical and mechanical level. Generally, NICs use PCI connections to the network. NICs can usually detect the type of network that they are connected to (e.g., 10 Base T, 100 Base T) and transmit at the appropriate rate.

A firmware connection for providing bootable flash system memory is usually also provided.

PRODUCT LIFE CYCLES

In order to choose a technology that will be useful over time, it is important to understand the timeframe in which the technology is available, as well as what update paths exist for migration. While software portability is paramount to providing a supportable infrastructure, awareness of where industry “is going” is also important. Various consortiums exist to promote products and technologies. Companies do not tie themselves to a single technology solution; they participate and create products in multiple technologies. For example, Compaq is strongly promoting PCI-X, while at the same time participating in the InfiniBand effort.

Chip end-of-life notice times can be very short. A product discontinuance notice may only provide a short time to procure the necessary hardware. For example, the Intel 440LX AGP chip set announcement of discontinuance occurred 8 June 2000, with total discontinuance reached on 8 December 2000—just six months later. Boards relying on this chip set can no longer be made and, therefore, will become obsolete when the stock of parts is depleted. Systems requiring these boards will also become obsolete.

PROCESSOR ROADMAPS

Processor roadmaps were investigated for the Intel, Alpha, Sparc (Sun), IBM, PowerPC, and Advanced Micro Devices (AMD) families of processors. Figure 2 shows the technology roadmaps for these vendors. (The appendix provides additional information.) It should be noted that this information changes over time as vendors update their roadmaps, new technologies become available, or vendors buy out other vendors. Value-line processing families and mobile market products, such as Intel Celeron, AMD Duron, or cell phone processors, are not included in this report.

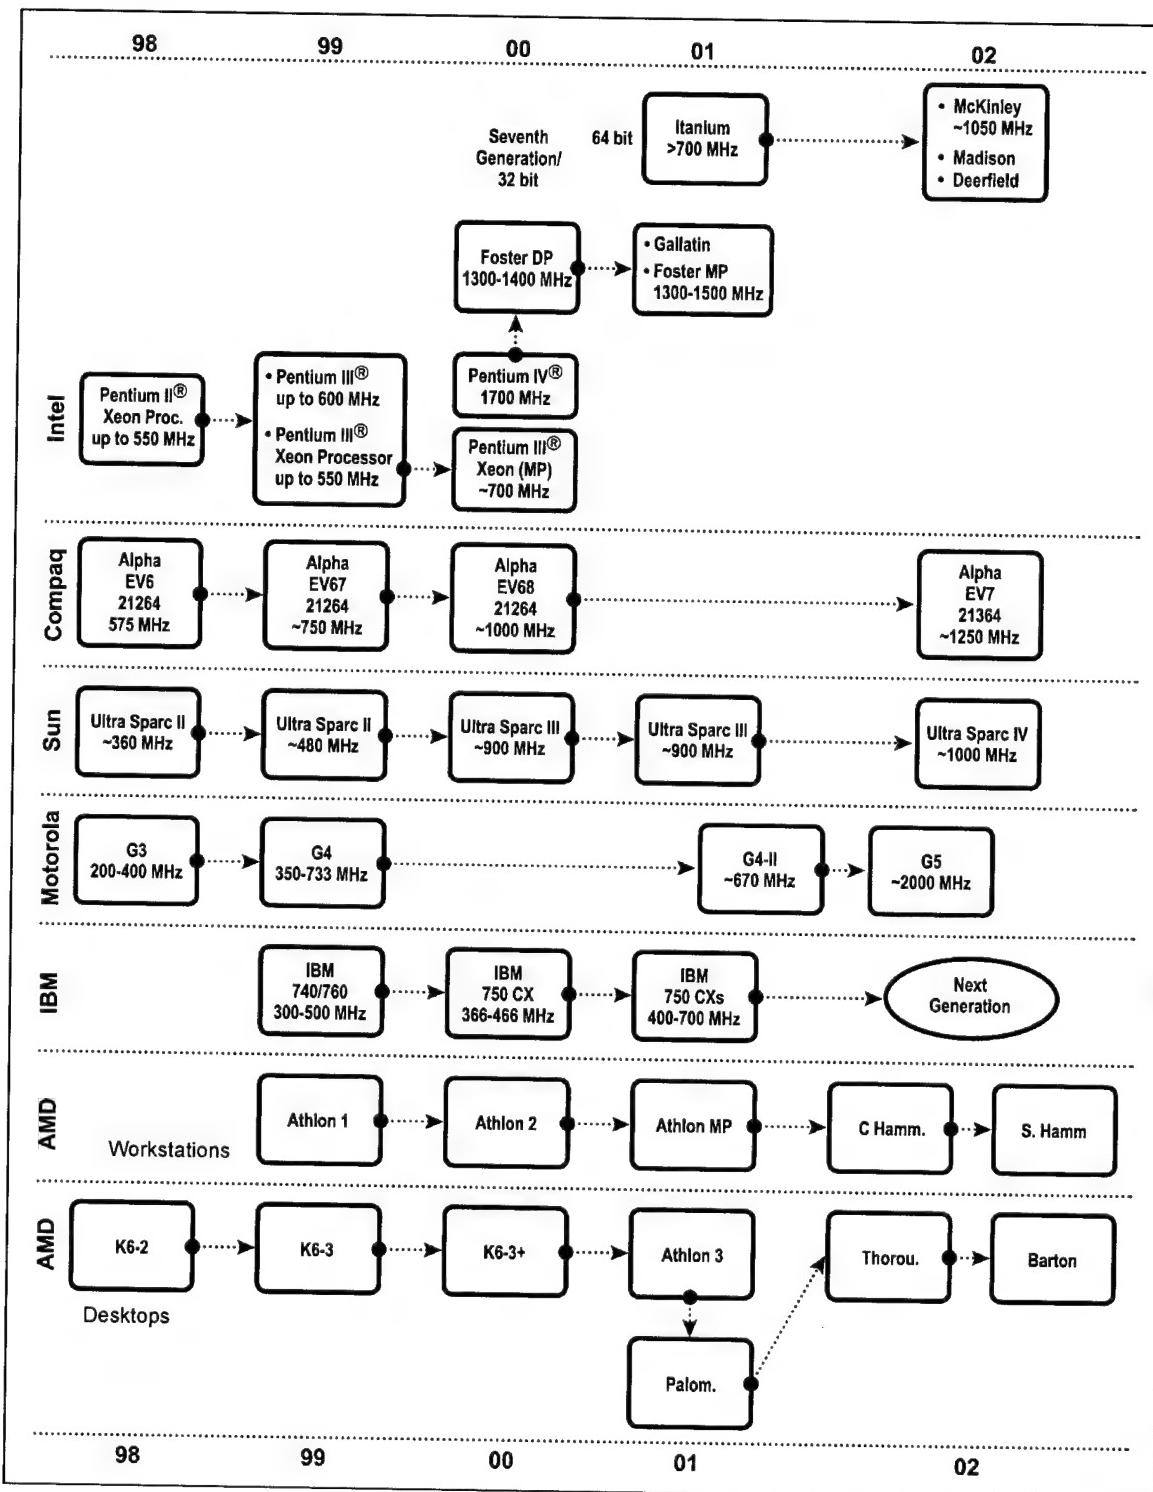


Figure 2. Technology Roadmap by Vendor

INTEL

Both Intel 32-bit and 64-bit architecture roadmaps were examined (reference 2). The 32-bit architecture (IA-32) will continue with the sixth-generation processors while adding a seventh-generation (Pentium IV). The 64-bit architectures will be provided by the Itanium family.

IA-32 Architecture

The sixth-generation Intel architecture reduces the chip size of the processor. In addition, new instructions have been added to support multimedia. Cache size has been increased, runs at the full-processor clock speed, and is integrated onto the chip. Front bus speed is increased from 100 MHz to 133 MHz.

The seventh-generation Intel architecture (Pentium IV) adds a 10-stage branch direction pipeline, which is capable of processing 100 instructions (at once) and 48 concurrent loads/stores. The L2 cache is integrated onto the chip; some also have L3 on the chip. More support is provided for multiprocessing. Out of order speculative execution (multiple instruction executed in parallel) is performed by hardware.

IA-64 Architecture

IA-64 is the first 64-bit architecture machine by Intel and is named Itanium. The IA-64 allows multiple instructions to be executed in parallel (compilers perform most of the decisions with some hardware support to determine what gets executed in parallel). The architecture used is known as Explicitly Parallel Instruction Computing (EPIC). Support is available for very large memory applications (tens of gigabytes). This need is expected in internet commerce and database applications.

ALPHA

The Alpha processor (reference 3) is a 64-bit processor. Digital Equipment Corporation (DEC) originally designed the Alpha family; Compaq has bought out DEC. The EV6, a third-generation (21264) Alpha, uses a smaller compatible metal oxide semiconductor (CMOS) die than second-generation chips. Speculative instruction fetch is used to increase performance. The EV7 (21364) focuses on symmetric multiprocessing (SMP) performance and scaling and integrates more functions onto the chip. The EV8 processors (21464) add simultaneous multithreading technology (independent threads issue multiple function calls/clock cycle) and increase processor utilization. In addition, more functions are integrated onto the chip. A new manufacturing process is used.

Compaq has recently transferred Alpha microprocessor and compiler technology, tools, and resources to Intel (reference 4). The technology will be incorporated into Intel Itanium II chips sometime in 2004. Compaq will use the Itanium II processor for all 64-bit workstations.*

SPARC

The UltraSparc II was designed for compute-intensive netcentric, multimedia applications (references 5 and 6). It is a 64-bit architecture. The UltraSparc III was designed to support aggressive real-world networking environments, such as e-commerce, large corporate intranets, high-capacity web servers and online transaction processing. It is designed for massively scalable (100s of processors) applications. The UltraSparc IV is similar to UltraSparc III with some small architectural changes. The UltraSparc V will be a new processor design. Sun maintains binary compatibility across families. Sun uses odd-numbered generations (I, III, V) for new architectures and even-numbered generations (II, IV) to introduce new process (chip) technology changes.

PowerPC

The PowerPC family was developed jointly by Motorola (references 7 and 8) and IBM (reference 9).

Motorola ‘G’ Family

The Motorola G3 has an L2 backside cache and an additional integer unit compatible with the G2. The G4 has an L2 backside cache on a chip as well as an L3 cache, AltiVec technology, a better FPU, and support for SMP applications. A G4-II processor is also expected. AltiVec allows the processor to perform digital signal processing, or customized computations, efficiently without the addition of further processors. The G5 will add a new bus and pipeline, be backward-compatible, and offer higher speed.

IBM

IBM uses a 64-bit data and 32-bit address processor. It uses copper technology. The IBM roadmap shows a path from copper technology to silicon to low-K dielectric processors. Speeds will go from 300 MHz in 1999 to 2+ GHz in 2001. Other additions include PCI-X support and an on-chip L2 cache. IBM has announced a new form of silicon called strained silicon that will be able to boost chip speeds by up to 35% because of better heat dissipation and less heat produced within the processor (reference 10). Strained silicon may be “productized” by 2003 (reference 11).

*At this time, Hewlett-Packard and Compaq are merging; the effect on technology roadmaps is unknown.

AMD

AMD makes processors for the personal and business markets focusing on desktop and notebook systems. It makes processors, embedded processors, and network and memory products. Such companies as Compaq, HP, and NEC-CI use AMD products (references 12 through 16).

TYPES OF MEMORY

Memory comes in two types of packaging, i.e., SIMM and DIMM. SIMM uses 32-bit-wide addressing while DIMM uses 64-bit addressing.

The most common type of computer memory is dynamic RAM (DRAM). DRAM generally uses one transistor and one capacitor to represent one bit. Synchronous DRAM uses a clock to synchronize the memory chip and the CPU and can run at a faster clock cycle. Static RAM (SRAM) does not lose memory during power losses; however, it runs at a slower clock cycle than other types of RAM. Flash memory chips do not lose memory contents during power loss and the memory can be erased and reprogrammed. Once reprogrammed, memory is retained until the chip is re-flashed. Flash memory is starting to replace EEPROMS due to lower cost and higher bit density (reference 17).

Synchronous DRAM (SDRAM) is a DIMM-packaged, 64-bit memory designed for Pentium III chips (reference 18).

Double data-rate (DDR) SDRAM is similar to SDRAM, but it provides twice the transfer rate as conventional SDRAM. A new type of memory called direct rambus (DRAM or RDRAM) introduces a new CMOS DRAM developed by the Rambus Corporation (reference 19). This memory module uses the standard DIMM form factor, but a different pin out on the connection. A two-byte-wide data channel is used, resulting in a peak data transfer rate of 1.6 Gbytes per second. Intel currently uses it in the Intel 820 chip set. It is available with Intel Pentium IV products (reference 20).

I/O INTERCONNECTION TECHNOLOGIES

Interconnection technologies are important in allowing data and instructions to be sent to and from a processor at a rate that keeps the processor running efficiently. Figure 3 shows the interconnection technologies available for current and near-future computing systems.

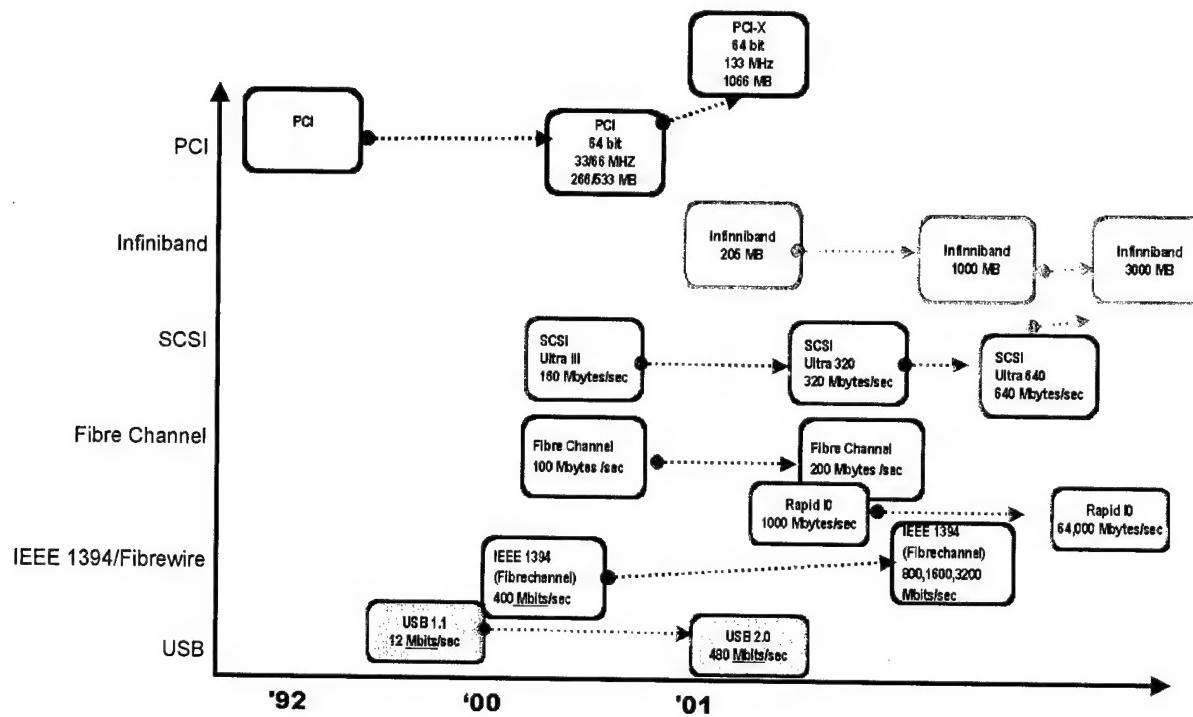


Figure 3. I/O Interconnection Technologies

BUS ARCHITECTURE

The standard architecture for passing data between the CPU and peripherals has traditionally been a bus (in particular, the PCI bus (figure 1)). The bus allows data to pass to and from peripherals, such as NICs, to networks and storage devices, such as disk drives. The bus typically interfaces to the memory controller of the system.

While buses form an important part of the computing infrastructure, technology packaging also plays an important role in deciding if a technology will “make it” in the market place — “...no interface enjoys wide success until vendors ship single-chip implementations” (reference 21). PCI is the standard in PC-computing markets. The versa module Eurocard (VME) bus has traditionally been used in military and aerospace environments; SCSI bus technology is primarily used in peripherals. Future emerging technologies include RACE and InfiniBand.

PCI and PCI-X

The PCI bus is the current PC bus standard used in PCs and workstations. It is used to pass data from devices (disk drives, networks, data storage devices) to and from the CPU and system memory of the workstation/computer. As applications become more I/O intensive, the need for a faster I/O channel becomes apparent. The PCI bus was first

marketed in 1992. The current PCI 2.2 spec is 64 bits and is coupled with either a 33.3-MHz or a 66.6-MHz bus. This allows a maximum peak bandwidth of 266 MB/s or 533 MB/s, respectively. It is believed that the current PCI 2.2 will not meet future needs, such as multi-port network interface cards with gigabit Ethernet (reference 22). PCI can only support two slots at the 66-MHz rate. In addition, other properties, such as error recovery hot swapping and reliability, are desired (reference 23).

CompactPCI is an industrial ruggedized variant of the PCI bus. CompactPCI technology is similar to desktop PCI, but with a different physical form factor. It was developed by the Peripheral Component Interconnect (PCI) Industrial Computers Manufacturers Group (ICMG). CompactPCI utilizes the Eurocard form factor popularized by the VME bus. CompactPCI is available in both 3U (100 mm by 160 mm) and 6U (160 mm by 233 mm) card sizes. CompactPCI has good shock and vibration characteristics (reference 24).

PCI-X is an evolutionary technology that upgrades PCI performance. It uses the 64-bit-wide bus of PCI with a 133.3-MHz system bus to provide a peak 1066 MB/s bandwidth. This bus is the joint work of Compaq, HP, and IBM. "This I/O bandwidth is needed for industry standard servers running enterprise applications such as Gigabit Ethernet, Fibre Channel, Ultra3 SCSI and Cluster Interconnects" (reference 25). The PCI-X spec is backward compatible; thus, PCI cards will work in PCI-X systems. In addition, at lower bus speeds, PCI-X can support more than two slots; four or more slots can be used at 66 MHz.

Currently, PCI-X is to be introduced in the server market. Eventually, PCI-X will be used in workstations and PCs. Since PCI technology revolves around the shared-bus topology, bottlenecks between competing users of the bus may still occur. InfiniBand (see below) is a follow-on technology that addresses this bottleneck.

VME

The VME bus specification was created in 1981 (reference 26). A VME bus board is either single (100 mm x 160 mm) or double in height (233 mm x 160 mm). A VME board is connected via a backplane, which can have up to 21 slots. The VME64 standard allows for 64-bit-block transfers. VME products have lower cooling requirements than conventional computers.

VME solutions are found in industrial controls, as well as military, aerospace, transportation, medical, and telecommunication applications. VME products include single-board computers, symmetric multiprocessors, communication products, memory, and interfaces. VME64 allows 80-MB/s transfers, while VME64x allows 160-MB/s transfers, and VME320 allows 320- to 500⁺-MB/s transfers. This performance has been available since 1997.

SCSI

SCSI (reference 27) is a peripheral bus technology that has been around for 20 years and is used for connecting devices (e.g., disk drives, CD-RW drives, and other mass-market low-cost devices) to a computer. The current version is the Ultra 3 SCSI with data transfer rates of 160 MB/s. Future generations include Ultra 320 and Ultra 640; Ultra 320 is expected to provide 320 MB/s data-rate transfer and Ultra 640 will provide 640 MB/s. Drive distance will limit connections to 12 m or less.

SWITCHED FABRIC CONNECTIONS

A switched fabric is an interface between any two devices. It provides a communication path without the need for the devices to know how they are connected. The fabric provides for multiple point-to-point connections to be made simultaneously by providing multiple input and output ports. The fabric moves the data between input and output. Important features include the ability to scale bandwidth requirements and provide high bandwidth connections. In addition, unlike bus technologies, switched-fabric performance does not decrease when nodes are added. Fibre Channel, RACE, RapidIO, Infiniband, PLX and Starfabric technologies will be discussed.

Fibre Channel

Fibre Channel (reference 27) is currently used in storage area networks (SANs) and redundant array of inexpensive disks (RAID) applications. Fibre Channel allows 1 Gb/s data transfers. SANs connect multiple storage systems together, usually at a centralized location, and allow access via multiple local area networks (LANs). SANs provide a scalable storage system and aid in data backup. Fibre Channel supports both SCSI protocol and internet protocol (IP). In addition, Fibre Channel can be run for fairly large distances and it has a large user-installed base in SANs. Fibre Channel allows transfers of 133 Mb/s, 266 Mb/s, 530 Mb/s, and 1 Gb/s on both electrical and optical media (reference 29). With overhead, the maximum data rate becomes 100 MB/s one way and 200 MB/s two way (reference 30). This data rate is theoretically available to each device in the connected fabric; a 2-GB version is expected shortly. Drive distances can be as long as 10 km using fiber-optic connections (FC). Current implementations use the FC-AL (arbitrated loop) to a host or device to connect to multiple hosts/devices (two loops are usually used in disk drives for failure).

Transmission control protocol (TCP)/IP is supported and is expected to replace SCSI over the next decade (reference 21). Drive distances can be as long as 10 km.

Race

Mercury's RACE architecture is a comprehensive, forward-looking heterogeneous, multicomputing "bus-less" architecture that is designed to specifically address the requirements of multichannel sensor-based I/O and processing in advanced signal processing, simulation, and communications applications. RACE's multicomputer interconnect, called RACEway, is a multi-node switching fabric that is designed as a scalable network of crossbar switches. At present, RACEway is implemented with six-port crossbar switches. Within each six-port crossbar, any of the six ports can be internally connected to any one or more of the remaining five ports at that crossbar. Up to three independent one-to-one port connections (i.e., internally connecting three unique pairs of ports) can be made simultaneously. Each crossbar port has a bandwidth of 160 MB/s (32 bits at 40 MHz). Hence, when the crossbar is used to make three independent port-to-port connections, the peak data transfer rate supported by each crossbar reaches 480 MB/s (reference 31).

RACE++ is the follow-on to the RACE switched-fabric architecture. RACE++ provides for up to 1-GB/s transfer bandwidth and is backward-compatible with RACE. Along with Mercury, third-party companies are also providing products. Markets include medical imaging and defense applications (reference 32).

RapidIO

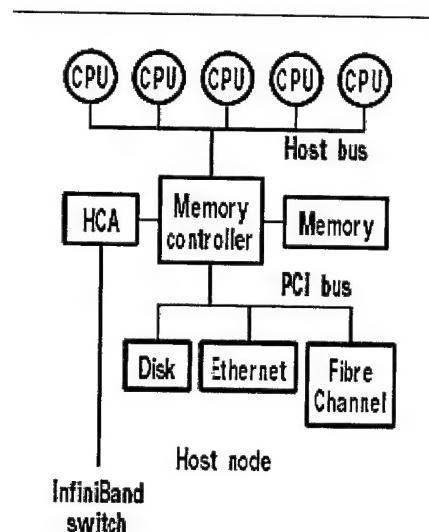
RapidIO is a switched-fabric architecture for connecting chips and boards within a system. It allows chip-to-chip and board-to-board communications and will be used in embedded systems, networking, wireless communications, and digital signal processing (DSP) applications. It is expected to provide performance to levels of 10 Gb/s. It may be used in integrated communications processors, host processors, networking, and digital signal processors (reference 33). RapidIO is used for connections within a chassis or box.

Currently, IBM, Alcatel, Cisco Systems, EMC Corporation, Ericsson, Lucent Technologies, Mercury Computer Systems, Motorola, and Nortel Networks comprise the consortium's steering committee with a total of 40 companies working on the specification (reference 34) for RapidIO.

InfiniBand

InfiniBand is a newly proposed I/O architecture. It is a future technology and may be the follow-on to PCI-X. InfiniBand will be used in multiprocessor clusters and remote I/O storage applications. It removes I/O communications from the CPU, allowing for more processing to be applied to the application (reference 35) versus using the CPU for communication (see figure 4). InfiniBand uses switched-fabric architecture to allow for multiple links and removal if the bus bottleneck. The InfiniBand Trade Association is developing InfiniBand. (Compaq, Dell, HP, IBM, Intel, Microsoft, Sun are steering committee members.) InfiniBand talks directly to the memory controller, bypassing the

bus. InfiniBand can talk to other computers, routers, or disk drives (reference 35), and can connect to system area networks and external networks. It uses the newer IP (IPv6 protocol). Data rates are provided for three different configurations. A one-wide link is expected to be 205 MB/s, a four-wide link is expected to provide 1 GB/s; and a 12-wide link is expected to provide 3 GB/s (in each direction). IBM, Compaq, Intel, HP, Sun, and Cisco are supporting InfiniBand. Products are expected sometime in 2001.



InfiniBand merges storage-area and system area networks and gets the bus out of the way.

Figure 4. PCI Plus Infiniband (Reference 35)

PLX

PLX (reference 36) is a new technology that provides a switched-fabric architecture that allows connection to legacy PCI devices. In addition, support for CompactPCI is provided. Drive lengths can be up to 4.6 m in length.

Starfabric

Starfabric was originally developed by Stargen (reference 37) and is now supported by Motorola, Bustronic, Natural MicroSystems, and Ziatech (now part of Intel). Starfabric provides switched-fabric architecture for CompactPCI applications. Starfabric will provide backward compatibility with existing systems using PCI and CompactPCI (reference 38). Bustronic Corporation, for example, is experimenting with a StarFabric hybrid backplane. “Systems with this backplane address next generation communications equipment requirements, such as trunk speed evolution from OC3 to OC12/OC48 and beyond, allowing 10-K to 100-K ports per single chassis” (reference 39). StarFabric can provide simultaneous support for packet, cell, and voice communications.

PERIPHERAL CONNECTIVITY

External devices need a connection path to the computer. Higher-speed devices, such as disk drives, are connected via a bus or fabric. Lower-speed peripherals, such as mice, scanners, keyboards, etc., are connected using a different standard. IEEE 1394 (firewire), universal serial bus (USB), parallel ports, serial ports, and the ISA bus provide peripheral connections.

IEEE 1394

IEEE 1394 (reference 40), or firewire, is a high-speed serial bus used for connecting relatively slow peripherals to computers. It is similar to USB, but provides a higher throughput and is more expensive. Firewire provides a data transfer rate of 400 Mb/s. Peripherals are “hot swappable,” i.e., they allow for connection and removal without powering down the unit or computer. Follow-on expectations include an increase in speeds from 800 to 1600, and, finally, to 3200 Mb/s.

USB

USB is used for connection of peripherals (reference 41) and provides a hot swappable bus with a transfer rate of 12 Mb/s. USB 2.0 is the follow-on with expected speeds of 480 Mb/s. USB is expected to replace serial and parallel port connections.

PARALLEL PORT/SERIAL PORT

Parallel ports (reference 42) have been used for connecting devices, such as printers, CDs, etc., to computers for some time. The maximum transfer rate is about 150 kB/s.

Serial ports also are used for connecting devices, such as modems, mice, and printers (although most printers are connected to a parallel port). A serial port, or interface, uses serial communication to transmit 1 b at a time. Data rates are 0.115 kb/s. Most serial ports on PCs conform to the RS-232C or RS-422 standards (reference 43).

ISA BUS

The bus architecture is used in the IBM PC/XT and PC/AT. It's often abbreviated as the ISA bus (pronounced as separate letters). The AT version of the bus is called the AT bus and has become a *de facto* industry standard. Starting in the early 1990s, ISA began to be replaced by the PCI local bus architecture. Most computers made today include both an AT bus for slower devices and a PCI bus for devices that need better bus performance.

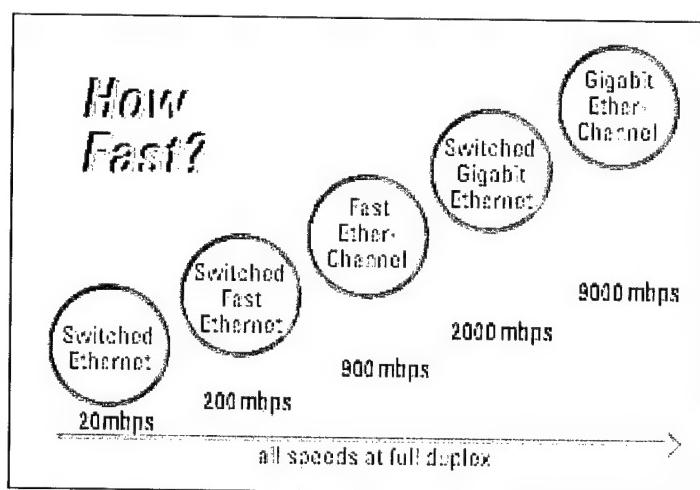
In 1993, Intel and Microsoft introduced a new version of the ISA specification called "Plug and Play ISA." Plug and Play ISA enables the operating system to configure expansion boards automatically so that users do not need to fiddle with DIP switches and jumpers (reference 44).

NETWORKS

Networks allow the interconnection of multiple machines and devices across small and large distances. Small distance (local) networks are called LANs while large distance connections are called wide-area networks (WANs). LANs typically inter-connect workstations, peripherals, terminals, and other devices in a single building or a relatively small geographic locale. LAN standards specify the cabling and signaling requirements at the physical and data link layer of the OSI reference model, embracing such communications technologies as fiber-distributed data interface (FDDI), Ethernet, and Token Ring.

ETHERNET TECHNOLOGY

Ethernet technology has been in use in industry since the 1970s. Ethernet uses a carrier sense multiple access/collision detection (CSMA/CD) protocol. Standard 10BaseT Ethernet has a 10 Mb/s transfer rate. There is a large installed user-base using Ethernet technology. Future increases in speeds across families of Ethernet are shown below in figure 5 (extracted from reference 45). Current work in Ethernet focuses on gigabit Ethernet technologies and beyond.



*Figure 5. Ethernet Technology: How Fast?
(Reference 45)*

Fast Ethernet

“Fast Ethernet (IEEE 802.3u) delivers 100-Mb/s bandwidth over category 5 unshielded twisted-pair (UTP) wire, or fiber-optic cable. This type of cable is commonly found in current network infrastructures. Like 10-Mb/s Ethernet, Fast Ethernet uses CSMA/CD network access method”(reference 46). Fast Ethernet is relatively cheap and provides a migration from standard Ethernet. Fast Ethernet is currently the primary LAN switching technology used in industry.

Gigabit Ethernet

Gigabit Ethernet uses the same frame format, frame size, and CSMA/CD protocol as Ethernet. A 1-Gb/s (1000-Mb/s) transfer rate (effective transfer of 200 to 400 Mb/s due to overhead) is possible. Gigabit Ethernet is used for network backbones and is compatible with installed Ethernet networks.

Gigabit Ethernet builds on top of the Ethernet protocol, but increases speed ten-fold over Fast Ethernet to 1000 Mb/s, or 1 Gb/s. Gigabit Ethernet provides high-bandwidth capacity for backbone designs while providing backward compatibility for installed media. Gigabit Ethernet can run over existing category 5 copper cabling; and, thus, it is an attractive choice to many companies with existing copper infrastructure. As the price for Gigabit Ethernet continues to decrease, Gigabit Ethernet will replace Fast Ethernet in LAN switching applications. Gigabit Ethernet will overtake Fast Ethernet by 2004 (reference 47).

10-Gigabit Ethernet (10GBASE-X)

The 10-Gigabit Ethernet is a follow-on to Gigabit Ethernet. The “10 Gigabit Ethernet uses the IEEE 802.3 Ethernet media access control (MAC) protocol, the IEEE 802.3 Ethernet frame format, and the IEEE 802.3 frame size. 10-Gigabit Ethernet is full duplex, just like full-duplex Fast Ethernet and Gigabit Ethernet; therefore, it has no inherent distance limitations. Because 10 Gigabit Ethernet is still Ethernet, it minimizes the user’s learning curve by maintaining the same management tools and architecture” (reference 48).

The 10-Gigabit Ethernet will be used in LAN, metropolitan area network (MAN), and WAN applications. The IEEE 802.3ae 10-Gigabit Ethernet Task Force controls these specifications. It is expected that pre-standard products will be available in 2001. The 10-Gigabit Ethernet specifications include plans to run over the Sonet high-speed network, thereby allowing for use in long-distance applications. Standards for 10-Gigabit Ethernet will not be finalized until 2002.

ASYNCHRONOUS TRANSFER MODE (ATM)

ATM is currently, the most widely used backbone technology. Today, ATM scales from T-1 to OC-48 at speeds that average 2.5 Gb/s in operation, 10 Gb/s in limited use, and spanning up to 40 Gb/s in trials (reference 49).

ATM is highly flexible, accommodating a wide range of traffic types, traffic rates, and communications applications. ATM interface standards exist for data rates as low as 1.544 Mb/s (DS1) and as high as 2.4 Gb/s (Sonet) (reference 50).

ATM interfaces can span the range from DS1/E1 rates to 622 Mb/s and beyond. Across this entire speed range, and across both the local and wide area, the common cell format and signaling protocols of ATM facilitate seamless internetworking and consistent service deployment. As network designers deploy LAN switches within wiring closets to alleviate workgroup congestion, there will be increasing demand for backbone and wide-area bandwidth to support the higher desktop speeds. ATM switching, spanning both the campus and the enterprise, will be ideal for this role, providing the orders of magnitude bandwidth increase over the desktop required for acceptable backbone performance. Over time, ATM may also extend directly to the desktop to allow new applications to use its other unique capabilities (reference 51). ATM provides guaranteed quality of service (QOS). In addition, all traffic types (e.g., voice, video, and data) can be provided.

SONET: HIGH-SPEED, FIBER-BASED TRANSMISSION MEDIUM

Sonet is a high-speed synchronous network specification developed by Bellcore and approved as an international standard since 1988. It is a fiber-based optical medium that has come into widespread use for data transport in broadband integrated services digital networks (BISDNs). This standard established a set of data rate and framing standards for data transmission using optical signals over fiber-optic cables.

The Sonet data rate and framing standards are designated as synchronous transport signal (STS-n) levels; the related Sonet optical signal standards are designated as optical carrier (OC-n) levels.

For the STS level, "n" represents the level at which the respective data rate is exactly "n" times the first level. For example, STS-1 has a defined data rate of 51.84 Mb/s; therefore, STS-3 is three times the data rate of STS-1, or $3 \times 51.84 = 155.52$ Mb/s. Similarly, STS-12 is $12 \times 51.84 = 622.08$ Mb/s, and so on.

Corresponding to each data rate and framing standard is an equivalent optical fiber standard. For example, the OC-1 fiber standard corresponds to STS-1, OC-3 corresponds to STS-3, and so on. The OC-n standard defines such items as fiber types and optical power levels (reference 50).

INTEGRATED PRODUCTS

Vendors are also providing integrated products that allow mixing and matching of network technologies. For example, the Cisco Catalyst 8500 switch provides both ATM and Gigabit Ethernet in a single chassis (reference 52).

STORAGE SOLUTIONS

Data storage capacity and transfer rates also play an important part in system performance. In addition, the ability to quickly recover from faults, damaged equipment, operator reboot, and the ability to upgrade deployed software is needed. Commercial storage capacity doubles every year and grows approximately 21% (in terms of dollars) (reference 53). Fibre Channel has a current large technology base in storage devices. Ethernet technologies are/will be used in storage solutions in the coming years. In the future, InfiniBand may also be used.

Hard drives are grouped together to form storage solutions. Direct attach storage systems and SANs form total storage solutions.

HARD DRIVES

Currently, both SCSI and integrated drive electronic (IDE) drives are in the market place. (See reference 54 for current/future products.)

SCSI Drives

U160-SCSI hard drives provide 160-MB/s transfer rates. SCSI controllers access several drives at once and share SCSI bus bandwidth. Generally SCSI drives are used in workstations.

IDE Drives

IDE hard disks integrate the electronics and firmware that previously existed on a separate controller card into the hard drive itself. A cache memory also exists to help speed up reading and writing to the disk. The IDE controller on the motherboard is a bus interface and a connector to the IDE cable is attached to the disk drive. Generally, a PC motherboard has two IDE interfaces that can support two IDE devices each. The term IDE is copyrighted by Western Digital.

The term ATA (AT attachment) is used by Maxtor, Quantum, and Seagate. ATA 66 has a top transfer rate of 66 MB. ATA 100, which offers 100-MB transfer, is now available (references 55 and 56).

Ultra DMA (or DMA/33 or ATA/33) is a protocol for ATA/IDE hard-disk drives. It is patented by Quantum. DMA/33 provides 33 megabyte per second transfer.

Enhanced IDE (EIDE) is a hardware interface connector used in IDE drives. IDE uses the system processor for most actions. This can slow down the system when the disk subsystem is under intense load. IDE drives are used primarily in desktop computing systems.

DIRECT ATTACH STORAGE (DAS)

Direct attach storage (DAS) is storage such as disk subsystems that attach directly to a server. Two general types in use are the redundant array of independent disks (RAID) and network attached storage (NAS).

RAID

RAID access and processing is generally through a single server, or CPU. (If the server goes down, so does access to RAID.) Most RAID drives are SCSI, although some now are being made with IDE drives. Figure 6 shows an example of a RAID drive by IBM (reference 57).

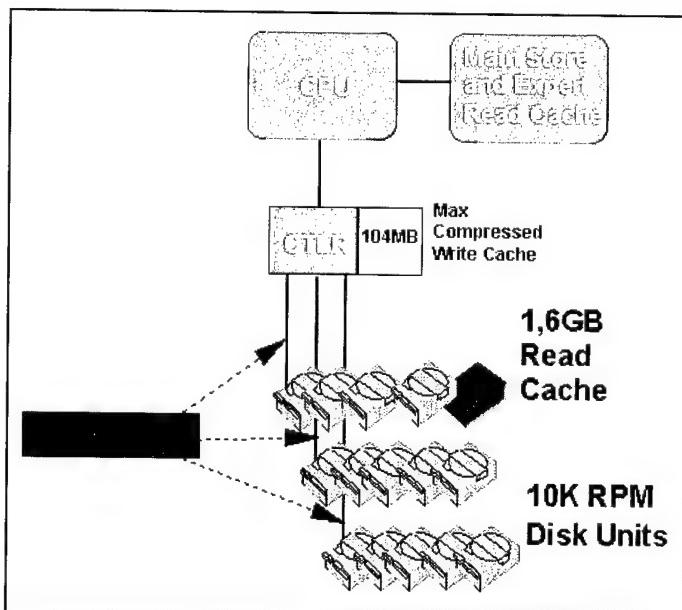


Figure 6. Example of RAID Drive (Reference 57)

NAS

In NAS (reference 58), a thin server (smaller computer) is used instead of a file server to access the storage. This “thin server” is attached to the disk subsystem, which frees up the processor that would have to be used to access the drives in a RAID system. NAS is less expensive than RAID drives. It attaches directly to the network. Servers can access a NAS if other servers are down. Figure 7 shows an example of a NAS.

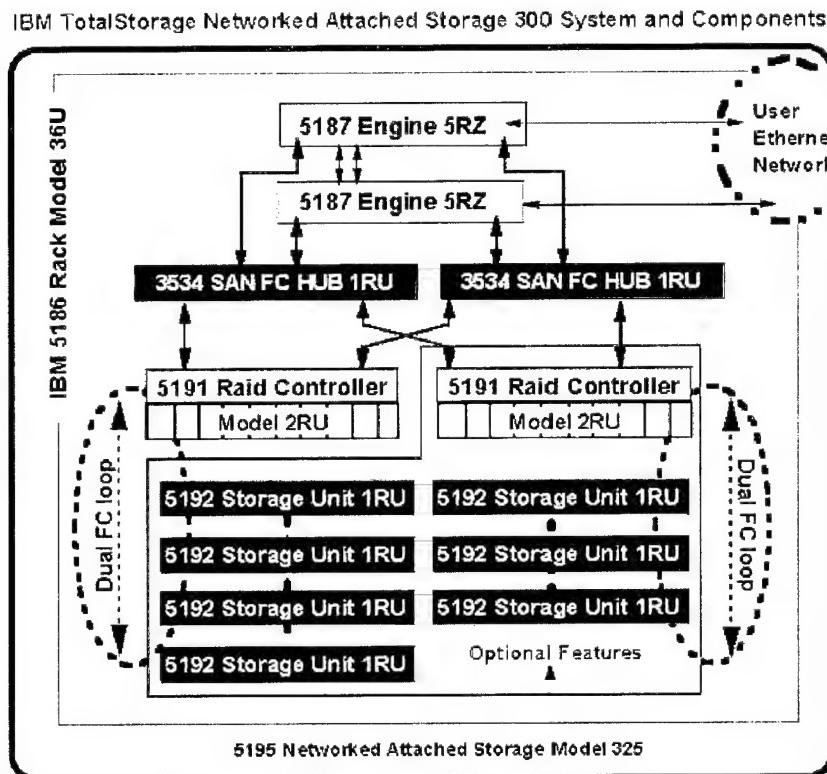


Figure 7. Example of a NAS Drive (Reference 58)

STORAGE AREA NETWORK (SAN)

SAN architecture (reference 59) links multiple storage subsystems (with their own LANs) together across a shared space. SANs connect multiple storage systems together (at several centralized locations) and allow access via multiple LANs. SANs can be used to provide a scalable storage system and to aid in data backup needs. In addition, since the SAN provides its own connectivity between devices, the use of SAN helps offload traffic from the LAN. The SAN can be constructed over great distances and are useful in disaster recovery. Figure 8 shows a typical SAN architecture.

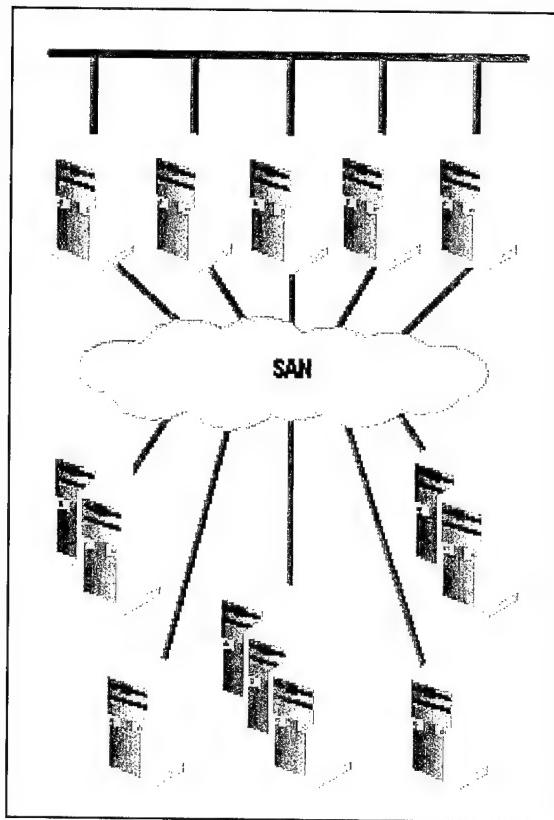


Figure 8. SAN Architecture (Reference 59)

SUMMARY

Computer hardware technology is constantly changing. While software architecture is very important, awareness of the direction of hardware technology change is needed for positioning future systems to take advantage of these changes and to ensure that these systems are maintainable in the future. Breakthroughs in technology, such as molecular computing, may lead to yet more explosive transformations. Whatever direction technology takes, the total computing architecture must be assessed in order to ensure needed future system performance, as well as to maintain an upgrade path for the computing infrastructure.

REFERENCES

1. "Architectural Support for Multiprocessing on Pentium II Xeon™ Processor-Based Workstations," <http://www.developer.intel.com/design/pentiumii/xeon/specupdt/243776.htm>.
2. "Itanium Processor Family," <http://developer.intel.com/design/IA-64/IA64RoadmapFall/sld001.htm>
3. "Computer Architecture," <http://www.comp.nus.edu.sg/~wuyinghu/academic/cs3220.html>.
4. "Compaq and Intel to Accelerate Enterprise Server Roadmaps," <http://www.compaq.com/newsroom/pr/2001/pr2001062501.html>.
5. "Microprocessor: UltraSparc Microprocessor Family," <http://www.sun.com/microelectronics/UltraSPARC/>.
6. "Sun Unveils Comprehensive Sparc Roadmap," <http://www.sun.com/smi/Press/sunflash/1998-09/sunflash.980901.1.html>.
7. "Motorola PowerPC Microprocessor Strategy," <http://e-www.motorola.com/collateral/PPCRMAPPDF.pdf>.
8. "High Performance, Low-Power 32-Bit PowerPC RISC Microprocessor," <http://e-www.motorola.com/brdata/pdfdb/docs/MPC7410fact.pdf>.
9. "IBM PowerPC Product Roadmap," http://www.chips.ibm.com/products/powerpc/rdmap/rd_map.pdf.
10. "IBM's Strained Breakthrough Image Page," <http://www.research.ibm.com/resources/press/strainedsilicon/>.
11. "IBM Research News," http://www.research.ibm.com/resources/news/20010608_strained_silicon.shtml.
12. "AMD Athlon 4 Processor," <http://www.amd.com>.
13. "AMD Processor Roadmap," <http://www.amd.com/us-en/Processors/Technical>.
14. "Athlon 4 Dual Processor," <http://www.geek.com/news/geeknews/2001june/bch20010601006152.htm>.
15. "AMD Athlon 4: The Palomino is Here," <http://www4.anandtech.com/showdoc.html?i=1469&p=3>.
16. "CHIPGEEK: Microchips, Chipsets and DRAM," <http://www.geek.com/procspec>.

17. "Transcend: Excellent High-End Memory Solution for Sun," <http://www.transcend.com.tw/english/>.
18. "A1 Electronics: SDRAM Memory Review," <http://www.a1-electronics.co.uk/Memory/SDRAM.html>.
19. "The Chip Connection Specialists," <http://www.rambus.com>.
20. "Intel Pentium 4 Processor: Performance Headroom for Tomorrow," <http://www.intel.com/eBusiness/products/desktop/p4p/wp012502.htm>.
21. "EDN Access, Table of Contents (1998 Archives)," <http://archives.e-insite.net/archives/ednmag/reg/1998/020298/03toc.asp>.
22. "Compaq Technology Brief: PCI-X An Evolution of the PCI Bus," <http://compaq.com/support/techpubs/whitepapers/tc990903tb.htm>.
23. "PCI-X Boosts Bus Bandwidth to 1 Gbps," <http://www.ednmag.com>, 11 May 2000.
24. "CompactPCI," <http://www.picmg.org/compactpci.stm>.
25. "PCI-X Industry Enablement Program," <http://www.compaq.com/products/servers/technology/pci-x-enablement.html>.
26. "Promoting Open Standards for the Bus and Board Industry," <http://www.vita.com/>.
27. "SCSI Trade Association," <http://www.scsita.org/>.
28. "About FCIA," <http://www.fibrechannel.com/>.
29. "High-Speed Interconnect Pages," <http://www1.cern.ch/HSI/fcs/spec/overview.htm#b1>.
30. "SAN Facts," http://www.netclips.co.kr/san/san_fc_faqs.pdf.
31. "Fitting Architecture to Application: Choosing Between SMP and RACE," http://www.mc.com/Technical_bulletins/mtb4smp-race/smp-v-race.html.
32. "RACE⁺⁺ Technology," <http://www.mc.com/plusplus/>.
33. "RapidIO Trade Association," <http://www.rapidio.org/home>.
34. "RapidIO News Room," <http://www.rapidio.org/news/rta010430.htm>.
35. "Network Computing," <http://www.networkcomputing.com/1205/1205ws22.html>.
36. "High-Speed Interconnect Silicon and Software for Networking and Communications," <http://wwwplxtech.com/>.

37. "Stargen, Company Facts," <http://www.stargen.com/>.
38. "Switched Fabric: A Stitch in Time," <http://www.ednmag.com>, 12 April 2001.
39. "Bustronic and StarGen Unveil New 21-Slot StarFabric Hybrid Backplane at Applied Computing Conference," <http://www.bustronic.com/release/twelve.html>.
40. "1394 Trade Association," <http://www.1394ta.org/>.
41. "High-Speed Certified USB," <http://www.usb.org/>.
42. "Warp Nine Engineering: Parallel Port Background," <http://www.fapo.com/porthist.htm>.
43. "Webopedia Search Term: Serial Port," http://www.webopedia.com/term/s/serial_port.html.
44. "Industry Standard Architecture Bus," http://www.zdwebopedia.com/TERM/I/Industry_Standard_Architecture_ISA_bus.html.
45. "Gigabit Ethernet Technologies," <http://www.cisco.com/warp/public/779/largeent/learn/technologies/gigeth/apps.html>.
46. "Designing Internets for Multimedia," <http://www.cisco.com/univercd/cc/td/doc/cisintwk/ndg4/nd2013.htm#xtocid1293530>.
47. "Gigabit Ethernet Migrates to LANs," <http://www.gcn.com/vol20>, Government Computer News, 23 July 2001.
48. Strategic Directions Moving the Decimal Point: An Introduction to 10 Gigabit Ethernet," http://www.cisco.com/warp/public/cc/techno/lnty/etty/ggetty/tech/10gig_wp.htm.
49. "About ATM Technology," <http://www.atmforum.com/pages/aboutatmtechfs1.html>.
50. "ATM Technology," <http://www.cisco.com/univercd/cc/td/doc/product/atm/l2020/2020r21/clicard/sysover/atmtech.html>.
51. "The Next-Generation ATM Switch: From Testbeds to Production Networks," http://www.cisco.com/warp/public/cc/so/neso/vvda/atm/ngatm_wp.htm.
52. "Catalyst 8500 Series Switches," <http://www.cisco.com/warp/public/cc/pd/si/casi/ca8500/>.
53. "Salmon Smith Barney Internet Technology Strategy: Data Storage," <http://www.SSBistrategy.com>.
54. "EDN Access," <http://www.ednmag.com/reg/1999/012199/02cs.html>.
55. "Hard Disk Drive Guide," <http://www.duxcw.com/digest/guides/hd/hd3.htm>.
56. "Seagate: Ultra ATA," http://www.seagate.com/support/kb/disc/ultra_ata100.html.

57. "IBM Integrated Storage: Hardware Architecture," <http://www-1.ibm.com/servers/eserver/iseries/hardware/storage/hwarchitecture.html>.
58. "IBM Total Storage Network Attached Storage 300 Delivers Superior Performance and Availability," <http://www1.ibmlink.ibm.com/cgi-bin/master?xh=ey38VY43RjpnLk1UsenGnF9332&request=atext&parms=mainPage.txt&xhi=&xfr=F> .
59. "Welcome to the Business Center," http://www.dell.com/us/en/biz/topics/power_ps1q00-sans.htm.

APPENDIX
PROCESSOR ROADMAPS FOR VARIOUS VENDORS

Table A-1. Summary of Intel Processors

IA-32 Architecture			
Processor	Speed	Date	Comments
Pentium III	600 MHz	1999	Concurrent single instruction multiple data (SIMD); 70 new (multimedia) instructions; increased memory bus utilization; 0.18-micron technology. Planned migration to 0.13-micron technology.
Pentium III (XEON)	550 MHz – 800 MHz	1999-2000	Multiprocessor version of Pentium III. Faster; bigger cache size. Coppermine variants, higher speed up to 800 MHz; smaller die; L2 cache on chip; different packaging capabilities. Cascade smaller die XEON; L2 on chip; replacement for XEON 700 MHz (form fit + bus); 0.18-micron technology.
Pentium IV	1700 MHz	End 2000-2001	Seventh-generation processor; 100-MHz system bus, 400-MHz data transfer (buffering) hyperpipeline; 100 instructions in pipe at same time, 48 simultaneous load/stores dual-processor and multiprocessor variants planned; 0.18-micron technology die shrink variants planned as well.
Foster	1300 MHz	End 2001	Pentium IV in a different packaging; MP version has significant changes; 0.18-micron technology; power 50-70 watts.
Gallatin			Based on Foster with smaller die size; performance systems chip; 0.13-micron technology.

Table A-1. Summary of Intel Processors (Cont'd)

IA-64 Architecture			
Processor	Speed	Date	Comments
Itanium	> 700 MHz (faster later)	2001	Explicitly parallel instruction computing (EPIC) architecture; more parallel data streams; L3 caching; four simultaneous extended-precision or eight simultaneous single-precision floating-points per clock cycle; can execute IA-32 (X86) instruction set; 0.18-micron technology.
McKinley	~1050 MHz	2001 (late)	Better floating-point performance over Itanium; 0.18-micron technology.
Deerfield/Madison		2002	Smaller die, larger L3 cache; Deerfield geared for price/performance; Madison geared for top performance; 0.13-micron technology.

Table A-2. Summary of Alpha Processors

Processor	Speed	Date	Comment
EV6	575 MHz	1998	Third-generation Alpha-upgrade to the Alpha 21264 processor; faster storage; 0.35-micron compatible metal oxide semiconductor (CMOS) with six layers of metal.
EV67	~750 MHz	1999	64-bit RISC architecture; used in Compaq Alphaservers, 1-14 processors; 0.25-micron CMOS with six layers of metal.
EV68	~1000 MHz	2000	Smaller die.
EV7	~1250 MHz	2001	Higher memory and cache bandwidth up to 256 processor configurations; on-chip L2 cache; focus on large-scale SMP applications; 0.18-micron CMOS with six layers of metal.
EV8	~1650 MHz ~2500 MHz	~2004	Smaller die; static and dynamic instruction level parallelism; higher data bandwidth to chip; L2 cache on processor; out of order instruction execution plus multithreading; 0.13 micron CMOS. (Note: This technology has been bought by INTEL.)

Table A-3. Summary of Sparc Processors

Processor	Speed	Date	Comment
Ultra Sparc II	250-480 MHz	1999	64-bit Sparc V9 pipeline architecture; smaller die than Sparc I; second generation processor; higher clock frequencies; 21 watts at 400 MHz; 0.25-micron CMOS.
Ultra Sparc III	600-900 MHz (up to 1500 on roadmap)	2000	New CPU core; second generation; 64-bit Sparc V9 pipeline architecture; bigger cache; bigger memory bandwidth; on-chip multiprocessor support; scalable to hundreds of processors; 70 watts at 750 MHz.; 0.18-micron CMOS.
Ultra Sparc IV	1000 MHz	~2002	Some architectural changes. Interconnects copper based
Ultra Sparc V	1500 MHz	~2003	New architecture pipeline. Interconnects copper based.

Table A-4. Summary of Motorola Processors

Processor	Speed	Date	Comment
G3	200-400 MHz	1998-1999	64-, 32-bit bus; backside L2 cache, second integer unit; low power; families 740, 745, 750, 755; 0.22-micron CMOS; five layers of metal; 5.8-8 watts at 300 MHz.
G4	350-733 MHz	1999-2001	64-bit bus; L2 on CPU; L3 cache AltiVec, 7-stage pipeline, better FPU; support for SMP; 0.15-micron CMOS; five layers of metal; 5-11.5 watts at 400 MHz New micro-architecture to be introduced; aster clock speeds; silicon-on-insulator (SOI) technology; 35% speed increase expected along with 65% power reduction. G4-II to be introduced; seven-stage instruction pipeline; two additional interface units (IU) added; initially 0.15-micron migrating to 0.13-micron.
G5	2000 MHz (expected)	2001	(Power PC 7500) New bus and pipeline; backward compatible; 32 bit and 64 bit available; 2 GHz speeds expected; 0.13-micron CMOS; copper interconnection.

Table A-5. Summary of IBM Processors

Processor	Speed	Date	Comment
740/750	300-500 MHz	1999	Uses copper technology; 64-bit data; 32-bit address; 0.22- and 0.25-micron CMOS; 3.7-6.5 watts.
750 CX	366-466 MHz	2000	256 K L2 cache; 64-bit data; 32-bit address; 0.18-micron CMOS; 6 watts at 600 MHz.
750 CXE	400-700 MHZ	2001	64-bit data; 32-bit address.

Table A-6. Summary of AMD Processors

AMD Desktop Market			
Processor	Speed	Date	Comment
K6-2	450-550MHz	1998	AMD value processor; 3D Now technology; 9.3 million transistors; 0.25-micron technology; uses C4 flip chip interconnection technology.
K6-3	~450 MHz	1999	3D Now technology, tri-level cache processor; can total 2368 kB of cache in L1, L2, L3. First processor with 100-MHz front-side bus
Athlon (K7) Generation 1	500 MHz	1999	(Information for all Athlon processors) Currently on seventh generation processor; 0.18-micron technology; full speed on die L2 cache (256K). New 266-MHz system bus technology (on more recent Athlon processors); enhanced 3D Now technology.
Athlon (Thunderbird Generation 2)/ Generation 3	1-1.4 GHz	2001 1 st Half	See Athlon Generation 1. 266-MHz front side bus.
Palomino (Athlon 4)	1500 – 1633 MHz	2001 2 nd Half	Advanced instruction set that accommodates new 760 MP chipset. On die L1 and L2, 37.5 million transistors. Draws 20% less power than Athlon generation 3 chips.
Thoroughbred (Generation 2 Athlon 4)	~2 GHz	2002 1 st Half	Expected. 0.13-micron technology, supports L1, L2, L3 cache; 266-MHz bus speed; enhanced 3D Now technology.
Barton (Generation 2 Thoroughbred)	~2 GHz	2002 2 nd Half	Enhanced Thoroughbred design; supports L1, L2, L3 cache; 266-MHz bus speeds; enhanced 3D Now technology; will support SOI technology.

Table A-6. Summary of AMD Processors (Cont'd)

AMD Server/Workstation Market			
Processor	Speed	Date	Comment
Athlon (K7) Generation 1	500 MHz	1999	Information for all Athlon processors: Currently on seventh-generation processor; 0.18-micron technology; full speed on die; L2 cache (256K); new 266-MHz system bus technology (on more recent Athlon processors); enhanced 3D Now technology.
Athlon (Thunderbird) Generation 2	~500 MHz	2000	See above.
Athlon MP (Mustang)	~1-1.2 GHz	2001	Multiprocessor, supports ~4-MB L2 cache; 266-MHz bus speed; increased power management over Athlon 3; enhanced 3D Now technology.
Clawhammer (Generation 1 K8)	~2 GHz	2002 2 nd Half	64-bit chip using x86-64 instruction set; 0.13-micron technology; small die size, <100 mm sq; used in 1-2 processor machines.
Sledgehammer (Enhanced Generation 1 K8)	~2 GHz	2002 2 nd Half??	Enhanced Clawhammer design; used in 4-8 processor machines with a larger L2 cache size.

INITIAL DISTRIBUTION LIST

Addressee	No. of Copies
Program Executive Office, Mine and Undersea Warfare (PMS 411: E. Plummer, H. Reichel, J. Feuillet, K. Michaud)	4
Naval Surface Warfare Center Division, Dahlgren (K. Eng, R. Stevenson, J. Allender)	3
University of Connecticut (A. Shukla)	1
Defense Technical Information Center	2
Center for Naval Analyses	1